

THE BROKER TIMES

COMPLIANCE

Australian compliance pack

AI WhatsApp Agent Engine · thebrokertimes.com.au

This is not legal advice. It is a starting framework written for a typical Australian broking business. Your licensee's policy overrides everything here. Get sign-off before this touches a real client.

The core principle: **an AI on your business number is you.** Every message it sends carries your name, your credit representative number and your obligations. There is no version of this where "the bot said it" is a defence.

1. Credit advice — the big one

Under the *National Consumer Credit Protection Act*, suggesting a credit product or assisting someone to apply is a regulated activity. Best Interests Duty applies to credit assistance provided to consumers.

An automated assistant cannot discharge Best Interests Duty. It cannot assess someone's circumstances, and it cannot form a reasonable view about what is in their interest.

So the agent must stay firmly on the *general information* side of the line:

Allowed — what your business does, what documents you typically need, how the process works, who you are, booking a call, general explanations of concepts like what LMI is or what an offset account does.

Not allowed — rates or rate ranges, borrowing capacity figures, product or lender recommendations, approval likelihood, "you should fix" or "variable is better right now", structure advice, anything about their specific numbers.

The supplied system prompt encodes this. If you loosen it, you are the one taking the risk.

Test it before every launch and after every prompt change. Guide 4, Step 2.

2. Spam Act 2003

The Spam Act covers commercial electronic messages, and WhatsApp messages count. Three requirements: consent, identification, and a functional unsubscribe.

Consent. Someone messaging you first is inferred consent to reply about that enquiry. That is the safe ground, and it is where this build lives — it only ever responds. Sending an unprompted WhatsApp to someone from your database is a different thing entirely, needs its own consent basis, and needs approved templates.

Identification. The agent identifies your business by name. Keep your business display name in Meta accurate.

Unsubscribe. The scenario honours STOP, UNSUBSCRIBE, OPTOUT, OPT OUT, OPT-OUT, REMOVE ME and END. On any of those it writes `opted_out = true` and never replies again. Test this yourself before launch — an unsubscribe that doesn't stick is the fastest way to a complaint.

Add other wordings to the opt-out keyword list in module 2 if you use them elsewhere.

3. Privacy Act 1988

You are now collecting personal information — names, phone numbers, and whatever people volunteer about their income, debts and circumstances — and storing it in Make's data stores.

Update your privacy policy to cover:

- That you use an automated assistant on WhatsApp
- What it collects
- Where it is stored, and that some of it is processed offshore (Make's servers, and your AI provider's)
- How long you keep it
- How to request access or deletion

The collection notice. APP 5 says you have to tell people what you are collecting and why, at or before the time you collect it — not later, and not only in a policy they never open. This build does it in two places, and both need your details in them:

1. The scenario's very first message (module 50) carries the AI disclosure and a link to your privacy policy. That is why module 50 has a `[PRIVACY POLICY LINK]` placeholder.
2. The system prompt makes the agent add one short line the first time money is discussed, naming your business and linking the policy again.

If you strip either of those out, you have an agent collecting income, debts and hardship disclosures from consumers with no notice at all. Leave them in.

Cross-border disclosure is the part brokers most often miss. APP 8 applies when personal information goes overseas. Both Make and your model provider are likely offshore. Your policy needs to say so.

Deletion requests. Someone can ask you to delete their data. You need to be able to find and delete their records in both `wa_contacts` and `wa_messages`. Know how to do it before someone asks.

Data minimisation. The system prompt instructs the agent never to ask for TFNs, dates of birth, bank details or ID over WhatsApp. Leave that rule in. WhatsApp is not a document collection channel.

4. AI disclosure

Not currently a specific statutory requirement for this use in Australia, but:

- ASIC has been consistently clear that licensees are responsible for AI they deploy
- Misleading someone into thinking they are talking to a person is, at minimum, poor conduct and arguably misleading under Australian Consumer Law
- It is trivially easy to disclose and costs you nothing

The scenario sends the disclosure itself, in the very first message, before the model is ever called. That is deliberate: it is fixed wording that cannot drift, it costs nothing in AI tokens, and it goes out even if the model call later fails. The system prompt then tells the agent not to repeat it, and to answer plainly if someone asks whether they are talking to a person.

Leave both halves in place.

5. Record keeping and Best Interests Duty

Every message in both directions is written to `wa_messages` with a timestamp. The model's running summary of the contact is written to `notes` in `wa_contacts`.

For this to be useful as part of a BID file note:

- Copy the transcript into your CRM against the client record when a conversation converts to an application
 - The transcript supports the file note; it is not the file note. Your own record of needs, objectives and reasoning still has to exist
 - Keep transcripts for at least as long as your licensee requires records to be kept
-

6. Vulnerable clients

The highest-risk failure mode is an automated system cheerfully qualifying someone in distress.

The system prompt escalates immediately on hardship, arrears, default, bankruptcy, divorce, illness and death. Do not remove those triggers.

Two things to add yourself:

- Decide what happens after escalation. "A human will call" must be true, and quickly. An escalation that nobody actions is worse than no escalation.
 - Consider adding your own triggers for wording your client base actually uses.
-

7. What to do when it gets something wrong

It will, eventually.

1. Turn the scenario off
2. Pull the exact transcript from `wa_messages`
3. Contact the client yourself, as a person
4. Tell your licensee — same day if it involved advice, a figure, or a vulnerable client
5. Fix the prompt with that specific case as an example
6. Re-run the full Guide 4 test set
7. Note what changed and when

Keep a change log of every prompt edit with a date. If you are ever asked what the system was doing on a given date, you will want it.

8. Before you launch — get these in writing

- ☐ Licensee or compliance manager has reviewed the system prompt
- ☐ They have seen real test transcripts, including refusals
- ☐ Privacy policy updated for AI, data storage and offshore disclosure
- ☐ Someone is accountable for actioning escalations, and knows it
- ☐ Transcript retention decided and documented
- ☐ A named person reads transcripts daily for the first week